

1  
2  
3 UNITED STATES PATENT APPLICATION  
4 FOR

5  
6 METHOD, SYSTEM AND COMPUTER PROGRAM PRODUCT FOR MONITORING  
7 OBJECTS IN AN IT NETWORK  
8  
9

10  
11 Inventors:

12  
13 Frank VOSSELER

14 Martin RECKER

1

2

## FIELD OF THE INVENTION

3

4

5

6

The present invention relates generally to the monitoring of an information technological (IT) network, and more particularly to a method, system and a computer program product for monitoring objects in an IT network.

7

8

## BACKGROUND OF THE INVENTION

9

10

11

12

13

14

15

Nowadays, as information systems become ubiquitous, and companies and organisations of all sectors become drastically dependent on their computing resources, the requirement for the availability of the hardware and software components of an IT network and of services based on it, (hereinafter all three are generally referred to as "objects") is increasing while the complexity of IT networks is growing.

16

17

There are monitoring systems available, which enable the availability and performance of objects within an IT network to be monitored and managed.

18

19

20

21

22

23

24

25

26

27

28

29

For example, Hewlett-Packard offers such a product under the name "OpenView VantagePoint". A personal computer, server, network interconnecting device or any system with a CPU is called a node. The nodes of an IT network monitored by such a monitoring system are called monitored nodes. On a monitored node, a program or process runs as a background job which monitors the occurrence of certain events (e.g. error messages) at the node and generates event-related messages according to rules which can be defined by a user. Such a program or process is called an "agent". An agent is not limited to passive monitoring, e.g. by collecting error messages. Rather, it can carry out active monitoring of hardware and processes. For example, an agent can periodically (e.g. every five minutes) send requests to a process (e.g. an Oracle process) to find out whether the process is still active. A response saying that the process is

1 no more active (or the absence of a response) may also constitute an "event".  
2 The messages generated by the agents are collected by a monitoring server  
3 which stores and processes them and routes the processing results to a  
4 monitoring console by means of which an IT administrator or operator can view  
5 the status and/or performance of the IT objects.

6 A monitoring system of that kind increases the availability of the IT objects  
7 under consideration since it enables a fault or failure of a component of the  
8 monitored network to be quickly detected so that a repair action or the like can  
9 immediately be started.

10 In such a known monitoring system, the operator is often overwhelmed with  
11 a large number of messages which are displayed at the monitoring console. Some  
12 of these messages are related to others, for example, when an application goes  
13 down and starts up again. Another example of related messages is the case in  
14 which a network router goes down so that all nodes beyond the router can no  
15 more be reached. As a consequence, the agent associated with the router will  
16 not only send messages indicating that the router is down, but will also send a  
17 larger number of messages that the nodes beyond the router are not available.  
18 Another fraction of the large number of messages sent to the monitoring console  
19 are messages which report similar or identical events, for example three  
20 messages reporting that a user switches to user root three times. Another  
21 fraction of messages report how a problem situation evolves with time, for  
22 example, how the amount of free disk space decreases or increases on a  
23 monitored node.

24 In some monitoring systems, messages can be correlated in order to  
25 suppress redundant messages, such as messages reporting identical events or the  
26 evolution of a problem situation. This is achieved by a kind of superordinate  
27 correlation analysis which is carried out by the monitoring server. However, the  
28 definition of this correlation analysis which has to be provided by the user, is  
29 rather complicated. Further, the correlation analysis requires considerable

1 computing and database access resources. A correlation of messages is not  
2 always possible since the available messages do not always contain the full  
3 information required for a correlation check.

4 Therefore, a monitoring system is desirable in which identical or related  
5 events can be more easily suppressed.

## 6 7 SUMMARY OF THE INVENTION 8

9 The invention provides a method for monitoring objects within an  
10 information technological (IT) network having monitored nodes in which  
11 monitoring-relevant events occur. The method comprises: generating event-  
12 related messages comprising a message key and a message relation key;  
13 comparing the message relation key with the message key of another message;  
14 and processing the other message depending on the result of the comparison.

15 According to another aspect, the invention provides a system for monitoring  
16 objects within an IT network having monitored nodes in which monitoring-  
17 relevant events occur, and a message processor. Agents are associated with the  
18 monitored nodes and generate event-related messages which comprise a  
19 message key and a message relation key. The message processor compares the  
20 message relation key with the message key of another message and processes  
21 the other message depending on the result of the comparison.

22 According to still another aspect, the invention is directed to a computer  
23 program product including program code for execution on a network having  
24 monitored nodes in which monitoring-relevant events occur. The computer  
25 program product comprises said program code for: generating event-related  
26 messages comprising a message key and a message relation key; comparing the  
27 message relation key with the message key of another message; and processing  
28 the other message depending on the result of the comparison.

29 According to yet another aspect, the invention is directed to a computer

1 program product including program code for execution on a network having  
2 monitored nodes in which monitoring-relevant events occur. The program code  
3 generates event-related messages comprising a message key and a message  
4 relation key for a comparison of the message relation key with the message key  
5 of another message. The message key and the message relation key are  
6 generated according to key patterns which can be defined on the basis of a set of  
7 pattern definition rules, and wherein both the message key pattern and the  
8 message relation key pattern are defined on the basis of the same set of pattern  
9 definition rules.

10 Other features are inherent in the disclosed system, computer program  
11 product and method or will become apparent to those skilled in the art from the  
12 following detailed description of embodiments and its accompanying drawings.

#### 14 DESCRIPTION OF THE DRAWINGS

15  
16 In the accompanying drawings:

17 Fig. 1 shows a high-level architecture diagram of a monitored IT  
18 network;

19 Fig. 2 shows an example of an event, an event pattern and a message  
20 key pattern, and a message key

21 Fig. 3 illustrates the generation of a message with a message key  
22 and a message relation key;

23 Fig. 4 illustrates how the message of Fig. 3 is compared with other  
24 messages

25 Fig. 5a,b show an exemplary view of message keys displayed in a message  
26 browser before and after the comparison carried out according to Fig. 4;

27 Fig. 6 illustrates the common usage of rules for message key patterns  
28 and message relation key patterns;

29 Fig. 7 illustrates that the message key pattern and the message relation

1 key pattern are defined at a common place in a computer program;

2 Fig. 8 is a flow-chart of a message correlation process carried out by an  
3 agent and a monitoring server.

#### 4 DESCRIPTION OF THE PREFERRED EMBODIMENTS

7 Fig. 1 shows a high-level architecture diagram of a preferred embodiment.  
8 Before proceeding further with the description, however, a few items of the  
9 preferred embodiments will be discussed.

10 In the preferred embodiments, objects of an IT network are automatically  
11 monitored as to their availability and performance by a monitoring system. Such  
12 monitored objects comprise hardware devices, software and services. A node is a  
13 network object such as a PC, server or any system with a CPU. A node is called  
14 a "monitored node" if it and/or applications or processes running on it are  
15 monitored by the monitoring system. Services are, for example, customer-based  
16 or user-oriented capabilities provided by one or more hardware or software  
17 components within a computing environment. For instance, services can be  
18 commercial applications (such as Oracle), Internet server applications (such as  
19 Microsoft Exchange), or the internal Windows NT 2000 services.

20 The monitoring may comprise passive monitoring (e.g. collecting error  
21 messages produced by the objects) or active monitoring (e.g. by periodically  
22 sending a request to the object and checking whether it responds and, if  
23 applicable, analysing the contents of the response). Besides pure monitoring  
24 tasks, in the preferred embodiments the monitoring system can also carry out  
25 management tasks, such as error correcting or fixing tasks, setting tasks and  
26 other network services control tasks.

27 An event is a (generally unsolicited) notification, such as an SNMP trap,  
28 CMIP notification or TL1 event, generated e.g. by a process in a monitored object  
29 or by a user action or by an agent. Typically, an event represents an error, a

1 fault, change in status or performance, threshold violation, or a problem in  
2 operation. For example, when a printer's paper tray is empty, the status of the  
3 printer changes. This change results in an event.

4 An agent is a program or process running on a remote device or computer  
5 system that responds to monitoring or management requests, performs  
6 monitoring or management operations and/or sends event notification. In the  
7 preferred embodiments, the agents are designed to run on the monitored nodes.  
8 Preferably, there is one agent for each monitored node. If several applications or  
9 processes run on the same monitored node, they will be preferably monitored by  
10 one and the same agent associated with this node. The agent is configured by a  
11 set of specifications and rules for each application or process to be monitored.  
12 The rules tell the agent what to look for and what to do when an event occurs  
13 (and, what events to trigger, if the agent carries out active monitoring). For  
14 example, according to particular rules, an agent filters events and generates  
15 messages which inform the monitoring server about the occurrence of certain  
16 events and/or the status and performance of the monitored application or  
17 process. In the preferred embodiments, the message processor is a part of the  
18 monitoring server. The monitoring server collects event and performance data,  
19 processes them and routes the messages and the results to a monitoring console  
20 (an user interface). In the preferred embodiments, services are monitored  
21 platform-independently. For example, different operating systems can be  
22 implemented on the various monitored nodes.

23 In the preferred embodiments, the generated event-related messages  
24 comprise what is called a message key. Preferably, a message key has attributes  
25 which characterise certain characteristics of the event that triggered that  
26 message, in particular those event characteristics which are interesting for  
27 monitoring purposes. For example, if the event is an entry in a logfile of a  
28 hardware device, the message key of the event includes only that part of the  
29 logfile entry which is considered relevant for the monitoring process. Thus, the



1 message key generation can represent a data filtering which removes irrelevant  
2 event information. However, which data are actually included in the message key  
3 can be freely chosen by the user through the definition of a message key  
4 generation rule which is called message key pattern.

5 Further, at least some of the generated messages comprise what is called a  
6 message relation key. The message relation key is an indication of which other  
7 messages (which are previously generated messages, simultaneously or  
8 subsequently generated messages) are related to the present message. The  
9 generation of the message relation key can also be defined by the user by means  
10 of a message relation key pattern.

11 The message relation key can be identical with the message key. In this  
12 case, only identical messages are considered as related to the present message.  
13 However, preferably the message relation key has at least one attribute which is  
14 a placeholder symbol (a wildcard). Such a wildcard leaves open the specific value  
15 of the data appearing at the position of the wildcard. Thus, a message relation  
16 key with a wildcard relates the present event to a class of other preferably  
17 previous events which are identical in the message relation key's attribute except  
18 the wildcard, but may differ from it in the attribute which is represented by the  
19 wildcard. For example, if a message key reports the fraction of available disk  
20 space of a particular hardware device, a message relation key with a wildcard at  
21 the available-disk-space-friction parameter relates the present message to all  
22 previous messages which reported the available disk space of that hardware  
23 device (irrespective of the concrete fraction value which was specified in those  
24 previous messages.

25 In the preferred embodiments, the message relation key of the present  
26 message is compared with the message keys of other, e.g. previous messages.  
27 Preferably, the comparison is a check as to whether the message relation key of  
28 the present message and the message key of the other message match with each  
29 other of course, taking into account possible wildcards in the sense explained



1 above.

2 The further processing of the other messages depends on the result of the  
3 comparison. In the preferred embodiments, the other message is discarded if it is  
4 found that its message key matches with the message relation key of the present  
5 message (of course, taking into account possible wildcards). Alternatively, rather  
6 than discarding previous related messages, the finding of a message correlation  
7 can be used for grouping messages together, weighting correlated messages in a  
8 particular way, displaying them in a particular way etc.

9 In the preferred embodiments, at least one monitoring agent is associated  
10 with a monitored node and generates the event-related messages comprising the  
11 message key. If the messages comprise also a message relation key, it is also  
12 generated by the agent. This is in contrast to common correlation methods in  
13 which relations between individual events (e.g. messages) are evaluated on a  
14 level above the generation of the individual events (e.g. at a monitoring server). In  
15 the present preferred embodiments, however, the correlation rule is entered at  
16 the lower level of the individual message generation and is transformed there into  
17 a message-specific message relation key. At the higher level (e.g. at the  
18 monitoring server), only a simple comparison between the message relation key  
19 and the other messages is carried out.

20 In the preferred embodiments, the message key and the message relation  
21 key are generated according to key patterns. A key pattern is a general rule  
22 specifying whether a message key is to be generated and which event-data are to  
23 be included into the message key or message relation key. Such key patterns can  
24 be defined by a user on the basis of a set of available pattern definition rules,  
25 which is, for example, similar to a script program which can be defined on the  
26 basis of a set of available script definition rules. Both the message key pattern  
27 and the message relation key pattern are defined on the basis of the same set of  
28 pattern definition rules i.e. on the same syntax. In contrast, in solutions in which  
29 the message key generation and the correlation analysis are performed on

1 different levels and are therefore logically different, the definition of key patterns  
2 and relations patterns would rely on different definition rules.

3 Further, in such solutions, the user would have to define the message key  
4 patterns and the correlation patterns at completely different places of the  
5 monitoring computer program. In contrast, in the preferred embodiments the  
6 message key pattern and the message relation key pattern can be defined by the  
7 user via a user interface at a common place of the monitoring computer program,  
8 since both patterns are used for the message key and message relation key  
9 generation by one and the same agent.

10 The preferred embodiments of the computer program product comprise  
11 digital program code which, for example, is stored on a computer-readable data  
12 carrier or is in the form of signals transmitted over a computer network. The  
13 preferred embodiments of the program code are written in an object-oriented  
14 programming language (e.g. Java or C++). The program code can be loaded (if  
15 needed, after compilation) and executed in a computer or in networked  
16 computers, e.g. a monitoring server networked with monitored servers.

17 Returning now to Fig. 1, it shows a high-level architecture diagram of a  
18 preferred embodiment of a monitoring system 1. The system 1 comprises  
19 monitored nodes 2, for example server computers on which applications 3, such  
20 as SAP<sup>TM</sup> 3a or Oracle<sup>TM</sup> 3b, run. The servers 2 are nodes of a monitored IT  
21 network, the network itself and the network interconnects are not shown in  
22 Fig. 1. The applications 3a, 3b may run on different platforms with different  
23 operating systems forming a heterogeneous network. A monitoring software  
24 component 4 (an "agent") is installed on each of the monitored servers 2 which  
25 runs automatically as a background task. The agents 4 receive event notifications  
26 and collect performance data from the monitored applications 3a, 3b and from  
27 hardware resources (server computers 2) used by them. They collect and  
28 evaluate these event notifications and performance data according to rules 5, 6.  
29 The rules comprise event patterns 5a, 6a, message key patterns 5b, 6b and

1 message relation key patterns 5c, 6c which are defined by a user via a graphical  
2 user interface (GUI) 7.

3 Depending on which events occur and what is indicated by the event  
4 patterns 5a, 6a and message key patterns 5b, 6b, the agents 4 filter the event  
5 data and generate monitoring messages 8', 8'', 8''' according to the message  
6 key patterns 5b, 6b, and send the monitoring messages 8', 8'', 8''' to a message  
7 processor, here a monitoring server 9. The messages 8 comprise a message key  
8 10b which is generated by the corresponding agent 4 according to what is  
9 prescribed by the corresponding message key pattern 5b, 6b. Messages 8 which  
10 are to be correlated with previous messages further comprise a message relation  
11 key 10c which is generated by the corresponding agent 4 according to what is  
12 prescribed by the corresponding message relation key pattern 5c, 6c. The  
13 monitoring server 9 stores the messages in a monitoring database 11, processes  
14 them and sends the messages and the processing results to a navigator display  
15 12 including a message browser 13. In the navigator display 12, the network and  
16 the services provided by it are visualised for the user in the form of a two-  
17 dimensional network and service map showing the status of the individual  
18 monitored services. In the message browser 13 the messages are displayed.

19 The monitoring server 9 comprises a comparator component 14 which  
20 compares the message relation key 10c of a newly incoming message 8 with the  
21 message keys 10b of all previous messages displayed at the message browser  
22 13. If it finds a match between the message relation key 10c and a message key  
23 10b of a previous message, it discards the previous message from being  
24 displayed at the message browser 13. However, the newly incoming message is  
25 sent to and displayed at the message browser 13 so that it effectively replaces  
26 the matching previous message. The network map in the navigator display 12 is  
27 automatically updated so that it reflects a status which corresponds to the newly  
28 incoming message rather than to the matching discarded message. In other  
29 embodiments, the comparison and discarding of messages can likewise be locally

1 performed at the message browser.

2 In some of the preferred embodiments, messages are only discarded from  
3 the message browser 13, but the database 11 keeps full record of all messages.  
4 In other embodiments, messages are also discarded from the database 11.

5 As will be explained in more detail below, the correlation of different  
6 messages can be freely defined by the particular choice of the message relation  
7 key pattern 5c, 6c and therefore, does not necessarily need to correlate identical  
8 messages (although it enables this to be done). Preferably, a separate mechanism  
9 for the suppression of duplicate events is provided. Since generally duplicates  
10 events originate from the same agent 4, the suppression of duplicate events can  
11 be carried out on the level of the agents 4. To this end, an agent 4 stores the last  
12 message it has generated and forwarded to the monitoring server 9. If the next  
13 message turns out to be identical with the stored previous message, it is not  
14 forwarded to a monitoring server 9. However, in other embodiments the agent 4  
15 forwards also duplicate messages to the monitoring server 9 for counting  
16 purposes. The duplicate messages are then identified and discarded by the  
17 monitoring server 9. In each case, the identification of duplicate events can be  
18 based on a simple identity check and does not require a comparison of a message  
19 relation key with the message keys of previous messages, as is the case with the  
20 general correlation mechanism which is explained in more detail below.

21 Fig. 2 illustrates the generation of a message key by means of an example.  
22 The exemplary event is a bad "Switch User" (SU) which has appeared on a node  
23 named "trex" from "frankv" to "root". This event shows up in a logfile as a  
24 syslog entry 21 as indicated in Fig. 2. Apart from the data "SU:-" (which means  
25 "bad SU"), "trex", "frankv" and "root", the logfile entry 21 carries a date and  
26 time indication, a terminal type indication ("ttyp3") and a "syslog" indication.

27 Rules 5 define whether and how a message key 10b is generated from the  
28 logfile entry 21. The rules 5 comprise an event pattern 5a and a message key  
29 pattern 5b if a message is to be correlated with other messages, the rules will

1 further comprise a message relation key pattern. This is explained in connection  
 2 with Fig. 3. The event pattern 5a has two functions: 1. It represents an IF clause  
 3 since it requires the logfile entry to be a "syslog" entry, and that a bad switch  
 4 user (i.e. "SU : -") has appeared. Only if this condition is met, a message key 10b  
 5 and, thus, a message 8 is generated; 2. It defines that the data appearing at the  
 6 places of the expressions `<. node >` `<*.from>` and `<*.to>` are assigned the  
 7 variable names "node", "from" and "to". The message key pattern 5b represents  
 8 a clause, i.e. it defines how a message key 10b is generated from the logfile  
 9 entry 21, provided of course that the condition defined by the IF clause is met.  
 10 In the example of Fig. 2, the pattern message key 5b comprises text "bad SU  
 11 from ... to ... on..." and indicates that the values of the variables "from", "to"  
 12 and "node" have to be inserted in the text. It does not include event data which  
 13 are considered irrelevant, such as the date and time indication, the syslog  
 14 indication and the terminal type indication. As a result, the message key 10b  
 15 "Bad SU from frankv to root on trex) is generated. Thus, the message key  
 16 generation includes filtering data (since a message key is only generated if the  
 17 logfile entry 21 matches the event pattern 5a) as well as condensing data, since  
 18 irrelevant event data are not included in the message key 10b.

19 In the example of Fig. 3, the event to be reported by a message is a  
 20 violation of a threshold on disk utilisation. In the present example, the threshold  
 21 level is set to 95%. Thus, an event occurs if the utilisation of a disk exceeds  
 22 95%. The message key pattern 6b comprises the text "disk\_util:/...:95". If  
 23 the threshold violation occurs in an object/opt on a node trex, the corresponding  
 24 variable data are inserted into the text and the generated message key 10b reads:  
 25 "disk\_util:/opt:trex:95".

26 An example of a message relation key pattern 6c is also shown in Fig. 3.  
 27 The purpose of the message relation key pattern 6c is that all previous disk-  
 28 utilisation-related messages referring to the same object and node shall be  
 29 correlated with the present message, irrespective of their threshold level.

1 Therefore, the message relation key pattern 6c is identical with the message key  
2 pattern 6b, apart from a wildcard (< \* >) in place of the particular threshold level  
3 "95". Consequently, the message relation key 10c which is generated according  
4 to the message relation key pattern 6c is identical with the message key 10b,  
5 apart from a wildcard replacing the particular threshold level. Thus, it reads for  
6 the exemplary event: disk\_util:/opt:trex:< \* >.

7 The message key 10b and the message relation key 10c are then sent from  
8 the agent 4 to the monitoring server 9. As illustrated in Fig. 4, the monitoring  
9 server's comparator component 14 compares the present message relation key  
10 10c with all previous message keys 10b',10b'' which are displayed in the  
11 message browser 13. In the example shown in Fig. 4, there is a previous disk-  
12 utilisation message key 10b' from node trex which indicates that the disk  
13 utilisation has been 75%. This message is now outdated, since the disk  
14 utilisation on this node is now above the threshold of 95%. The message key  
15 10b' matches with the message relation key 10c of the present message. In  
16 particular, the wildcard in the message relation key 10c is considered as  
17 matching with the disk-utilisation value "75" of the message key 10b'. Owing to  
18 the positive determination of the matching test, the message with the message  
19 key 10b' is discarded from display. In contrast, another message 10b'' is also  
20 shown which does not match with the message relation key 10c (since it refers  
21 to node "frankv" rather than "trex"). Since the matching is determined to be  
22 negative in this case, the message with the message key 10b'' is not discarded  
23 from display.

24 Fig. 5a shows what is displayed by the message browser 13 before the  
25 comparison of Fig. 4 is carried out. The display shows two message keys which  
26 indicate that the disk utilisation is 75% on node trex and 95% on node frankv.  
27 After the comparison in Fig. 4 has been carried out, the first-mentioned message  
28 key has been replaced by the present message key 10b which indicates that the  
29 disk utilisation on node trex is 95%. Thus, the outdated message 10b' is no



1 longer displayed as soon as a more recent message is received from the  
2 corresponding agent 4.

3 As already indicated in Figs. 3 and 4, the message key patterns and the  
4 message relation key patterns use the same rule definition syntax. This is further  
5 illustrated by Fig. 6 which shows that both message key patterns 5b and  
6 message relation key patterns 5c are defined on the basis of the same set 22 of  
7 rules and rule elements. This fact is very advantageous since it discharges the  
8 user from the need to learn and use two different rule definition languages and,  
9 thereby, facilitates the set-up procedure and operation of the monitoring system.

10 Owing to the fact that both the message key and the message relation key  
11 are generated in response to the occurrence of an event by one and the same  
12 agent, and that both are based on the same rule definition syntax, the message  
13 key pattern 5b and the message relation key pattern 5c can be defined by a user  
14 via a common user interface 23 at a common place 24 of a monitoring computer  
15 program 25, as is illustrated in Fig. 7. The "common place" may be, for example,  
16 a common class or routine of the program.

17 Fig. 8 illustrates the message correlation method carried out by the agents 4  
18 and the monitoring server 9. In step S1 an event occurs. In step S2 the agent 4  
19 ascertains whether the event matches with one of the event patterns 5a  
20 comprised in the rules 5. If the answer is positive, in step S3 a message key 10b  
21 is generated according to the corresponding message key pattern 5b. If the rules  
22 5 comprise also a message relation key pattern 5c, in step S4 a message relation  
23 key 10c is generated, according to the message relation key pattern 5c. The  
24 steps S2 to S4 are carried out by the agent associated with the monitored node 2  
25 on which the event occurred. Then, in step S5, the message 8 with the message  
26 key pattern 5b and the message relation key pattern 5c is sent to the monitoring  
27 server 9. In step S6 the monitoring server 9 ascertains whether the message  
28 relation key 10c matches with the message keys 10b of the previous messages.  
29 If it finds a matching previous message it discards it in step S7 from display in



1 the message browser 13. In step S8 the monitoring server 9 initiates the display  
2 of the present message 8 in the message browser 13 so that it appears as if the  
3 matching previous message were replaced by the present message 8. The steps  
4 S6 to S8 are carried out by the monitoring server 9 in conjunction with the  
5 message browser 13. As soon as another event occurs, the sequence of steps  
6 S1 to S8 is carried out again, thus forming a quasi-endless loop.

7 Thus, a general purpose of the disclosed embodiments is to provide an  
8 improved method, computer system and computer program product for  
9 monitoring services in an IT network. In particular, the disclosed embodiments  
10 enable a definition of the correlation analysis by the user. The computing and  
11 database access resources needed for the correlation analysis are reduced. A  
12 correlation of messages is even then possible if the available messages do not  
13 contain the full information required for a correlation check. Thus, the disclosed  
14 embodiments allow an easier suppression of identical or related events.

15 All publications and existing systems mentioned in this specification are  
16 herein incorporated by reference.

17 Although certain systems, methods and products constructed in accordance  
18 with the teachings of the invention have been described herein, the scope of  
19 coverage of this patent is not limited thereto. On the contrary, this patent covers  
20 all embodiments of the teachings of the invention fairly falling within the scope of  
21 the appended claims either literally or under the doctrine of equivalents.